

# THE BIDIRECTIONAL LUNAR LOGISTICS NETWORK

*A persistent architecture for a permanent lunar base*

## Author:

Ernesto Gabriel Arias Fernandez

Independent Researcher, Lunar Logistics Architecture

redlunar.gabrielarias@gmail.com

---

## Abstract

Current lunar architectures remain anchored in the paradigm of the isolated mission: a giant rocket, a heroic lander, a predictable end. This document demonstrates, step by step, that this approach is inadequate for the goal of "staying." Starting from the fundamental principles of the Tsiolkovsky equation and  $\Delta V$  segmentation, the reader is led to discover a structurally different alternative: a bidirectional logistics network between Earth and the Moon.

The network consists of four nodes (LEO, TLI, LLO, SUP) and three types of vehicles (expendable cryogenic stage, orbital tug, reusable hypergolic lander). Its central element is the surface network: a distributed propellant depot formed by the landers themselves, interconnected, sharing energy, heat, and propellant.

Network simulation shows a sustained throughput of ~240 tons per year to the lunar surface with 12 annual launches, an efficiency 3 to 7 times higher than current architectures (SpaceX HLS, Blue Origin). All the necessary technology exists today. The only condition is to change the mental framework: from the epic mission to persistent infrastructure.

**Keywords:** lunar logistics network, permanent lunar base, hypergolic architecture,  $\Delta V$  segmentation, orbital refueling, surface network, space logistics.

---

## 1. Introduction

We are living through a historic window of opportunity in space exploration. For the first time, three factors coincide that have never occurred together [1][2][3][4][5]:

- A new space race focused on the Moon and permanent presence.
- Giant reusable rockets, capable of launching tens or hundreds of tons at low cost.
- Mature, robust, proven technologies that allow operating in space with unprecedented reliability.

This combination changes everything. For the first time, building lunar infrastructure ceases to be science fiction and becomes a realistic, economically viable, and technologically accessible project.

However, we continue to approach the problem with a mental framework inherited from the Apollo Program: isolated missions, heroic vehicles, unique flight profiles, landers that go up and down once and then die of success.

That approach no longer works.

The Tsiolkovsky equation —perfect for describing a single burn— does not describe a system. It does not account for availability, cycles, orbital permanence, structural penalties, or tons per year. That is why the "cryo vs hyper" debate is misframed: it compares engines, when what we need to compare are architectures.

Space is not moved by  $\Delta V$ . Space is moved by logistics.

And if we want to stay on the Moon —not just visit it— we need to change the fundamental question:

Not "which engine is better?", but "which system moves more mass, more times, with more reliability?"

The answer is clear: we do not need a perfect lander, nor a perfect rocket, nor a perfect mission.

*We need a network.*

A bidirectional logistics network between Earth and the Moon. A network that moves people, cargo, fuel, water, spare parts, tools, and products. A network that works every day, not every two years. A network that turns the Moon into an accessible place, not an exotic destination.

---

## **2. Mission Architectures vs. Logistics Networks**

Contemporary space architectures have mostly been developed under a mission paradigm. In this approach, each operation is conceived as an isolated event, individually optimized to maximize performance within a closed set of constraints. This model was extraordinarily effective during the initial phase of space exploration, when objectives were punctual, resources limited, and repetition was not a requirement.

But this same approach reveals deep limitations when evaluated from the perspective of sustained operations.

In a mission architecture:

- Each flight is independent.
- Infrastructure is minimal or nonexistent.
- Mass flow is discontinuous.
- Optimization focuses on a single event.

The result is highly specialized systems, energy-efficient, but difficult to scale, with low operational frequency and persistently high costs per kilogram.

In contrast, a persistent logistics network introduces a fundamental conceptual shift. The system ceases to be defined by isolated events and becomes defined by flow continuity.

In this framework:

- Permanent infrastructure,
- Operational frequency,
- And the existence of nodes and intermediate storage

become the dominant elements. Performance no longer depends on optimizing a single component, but on the interaction among all of them.

This change completely alters the relevant variables.

In mission architectures, specific impulse and energy efficiency are critical because each kilogram must be transported in a single event with no possibility of redistribution. In logistics networks, however, the existence of depots, repeated cycles, and stable routes redistributes energy constraints throughout the system. Sensitivity to Isp decreases, while variables such as throughput, operational reliability, and iteration capacity gain greater weight.

Risk management also changes. Mission architectures concentrate risk in single, highly complex events: a failure implies total loss of the objective. In a logistics network, risk is distributed across multiple operations. Individual failures can be absorbed without compromising system continuity, as long as the infrastructure remains intact.

Furthermore, logistics networks facilitate technological evolution. While a mission architecture remains anchored to initial decisions, a persistent network allows new technologies to be incorporated progressively, without redesigning the entire system. Infrastructure acts as an integration platform.

From an economic perspective, the difference is even more marked. Mission architectures present high costs per unit mass due to low frequency and absence of

economies of scale. In logistics networks, increased throughput and infrastructure reuse reduce cost per kilogram as a natural consequence of continuous operation.

For all these reasons, moving from mission architectures to logistics networks is not an incremental optimization: it is a paradigm shift.

It is not about designing the best possible mission. It is about building a system capable of sustaining activity over time.

And this distinction is fundamental:

Exploration can be achieved through missions. The utilization of space requires networks.

---

### **3. How Can We Build the Bidirectional Earth–Moon Logistics Network?**

If we want the main objective of today's competing lunar programs —this new 21st-century "Space Race," which according to their own declarations is the establishment of a Permanent Lunar Base at the Moon's South Pole— to stop being an aspirational dream and become a tangible reality, we have already established the need to build the Network. We have also said that we have the existing technology and the necessary level of maturity to do so. If we change the paradigm that has misdirected our efforts —the focus on isolated missions— and listen to what the laws of physics, mathematics, orbital mechanics, the Tsiolkovsky equation, and existing technology have to tell us, we can achieve it.

So, join me in discovering this Bidirectional Logistics Network.

---

### **4. The Topological Basis of the Network**

Every network consists of nodes. We must therefore choose ours. The most obvious and inevitable, to the point that it requires no explanation, is our initial node: the Earth's surface, more specifically, the launch pad where our journey begins. Let us call it Node\_Earth.

The rocket we will use to climb our first delta-V mountain (9,500 m/s) to the second node can be any, because the network must be conceptually independent of specific hardware. However, the existence of reusable launchers with enormous payload capacity to LEO is an essential part of the window of opportunity we are experiencing. To simplify understanding and the development of the next steps in construction, we will adopt Starship. But any other choice would be equally valid.

According to specifications that Elon Musk himself has repeatedly shared on social media, this vehicle will have a capacity in the range of 100 to 200 tons [3]. For our example, we will assume that Starship delivers 150 tons of payload to LEO. But we repeat, the network does not depend on a specific payload to LEO quantity; what we will develop in this example for 150 tons applies indistinctly to 100 tons, 55 tons, or 200 tons. Here we have, then, the second node of our network: Low Earth Orbit, LEO, which we name Node\_LEO.

At this point, we must choose the next node. The logical options in practice are two:

- **Low Lunar Orbit (LLO):** The one used by the Apollo missions. A valid alternative for missions and lunar bases near the lunar equator.
- **Near-Rectilinear Halo Orbit (NRHO):** A valid alternative for missions and lunar bases near the lunar poles.

It is not the purpose of this paper to argue the advantages and disadvantages of these orbits. Their specific choice depends on a series of factors beyond the scope of this text. But we can take advantage of the fact that their delta-V requirements are similar. From the two possible nodes, let us choose for our example Low Lunar Orbit (LLO), although everything we develop from here on is also applicable to NRHO.

Thus, our third node will be Node\_LLO.

The fourth node of this network will necessarily be the lunar surface, Node\_SUP.

The main consequence of our node choice is that it introduces a natural segmentation into the delta-V required for lunar logistics operations. Although there are small variations in delta-V between nodes when calculating a real trajectory for each specific mission, depending on countless factors (orbits and exact positions of vehicles, trajectory optimization and corrections, etc.), we will establish typical and representative delta-V values for each leg:

| Leg                               | $\Delta v$                                           |
|-----------------------------------|------------------------------------------------------|
| Node_Earth $\rightarrow$ Node_LEO | $\approx 9,500$ m/s                                  |
| Node_LEO $\rightarrow$ Node_LLO   | $\approx 3,100$ m/s (TLI) + 900 m/s (LOI to LLO)     |
| Node_LLO $\rightarrow$ Node_SUP   | $\approx 1,900$ m/s                                  |
| Node_SUP $\rightarrow$ Node_LLO   | $\approx 2,000$ m/s                                  |
| Node_LLO $\rightarrow$ Node_Earth | $\approx 900$ m/s (TEI)                              |
| Node_LLO $\rightarrow$ Node_LEO   | $\approx 900$ m/s (TEI) + 3,100 m/s (braking to LEO) |

This allows us to understand the formidable problem faced by the designers of the Artemis mission Landers, an architecture based on individual missions. The Artemis architecture implies that each Lander must perform the sum of several of these legs by itself, dragging its own dry mass along with the propellant mass for subsequent legs. This leaves little margin for "payload."

**The problem of mission architectures: the lander must do EVERYTHING.**

In Artemis, each Lander must perform the sum of several of these legs. Let us see it in numbers:

Typical Artemis lander profile (LEO → Surface → NRHO):

- LEO → TLI:  $\Delta v \approx 3,100$  m/s
- TLI → NRHO:  $\Delta v \approx 500$  m/s
- NRHO → LLO:  $\Delta v \approx 800$  m/s
- LLO → Surface:  $\Delta v \approx 1,900$  m/s
- Surface → NRHO:  $\Delta v \approx 2,800$  m/s

**Total accumulated:  $\approx 9,100$  m/s [14]**

This accumulated  $\Delta V$  is a true challenge, as it lies within the regime where the Tsiolkovsky equation fully manifests its logarithmic nature, penalizing available payload and far from the low  $\Delta V$  regimes, where the logarithmic curve remains fairly flat. This contradicts the intuition of the brilliant Russian [15], who was perhaps the first to suggest that the correct way to handle high delta-V situations. Tsiolkovsky was the first to warn that large  $\Delta V$  must be divided into small segments, creating the concept of the multistage rocket that we all know and thanks to which we have been able to reach space.

SpaceX's proposal seeks to overcome the challenge through brute force, resorting to massive refueling of propellants at Node\_LEO, in low Earth orbit, by deploying a true fleet of tanker Starships to replenish the lunar Starship. According to NASA, between 15 and 20 launches [13]. A true feat requiring perfect synchronization of these sequential launches to minimize propellant boil-off, active management to reduce it, and the transfer of cryogenic propellants, both topics we do not yet handle and remain in the realm of future technologies.

Blue Origin's proposal faces the challenge with its bet on a high-performance hydrolox engine, the BE-7 [4], a jewel of advanced engineering, dividing the  $\Delta V$  into two legs using a Transporter. Although its proposal is not yet well defined, it would probably need 6 or 7 flights of the NG 9x4. It is an elegant approach, but equally conditioned by the need for a single system to complete a huge accumulated  $\Delta V$ .

All this demonstrates the extraordinary capacity of our scientists and engineers. But it also demonstrates something deeper:

They are fighting against the limitations of a mission architecture. Not against the limitations of physics.

When we force a single vehicle to perform all legs, we are ignoring the fundamental lessons of Tsiolkovsky and logistics:

- Large  $\Delta V$  must be divided.
  - Payloads must be transferred.
  - Vehicles must specialize.
  - And the system must operate as a network, not as an isolated mission.
- 

## 5. Building the Network

We said we would leave the first leg to Starship. So, let's go to the Moon.

### Node\_LEO to Node\_SUP:

- LEO  $\rightarrow$  TLI:  $\Delta v \approx 3,100$  m/s
- TLI  $\rightarrow$  LOI (LLO):  $\Delta v \approx 900$  m/s
- LLO  $\rightarrow$  SUP:  $\Delta v \approx 1,900$  m/s

### From LEO to TLI:

We are in LEO with a total mass  $M_t = 150$  T that Starship has just delivered. To enter TLI, we will use a Cryogenic Upper Stage. This is not a whim: it is what Tsiolkovsky dictates. With  $\Delta v \approx 3,100$  m/s we enter the optimal regime for hydrolox.

This is the most significant and demanding leap of the entire network, so we will use the best of our current technology. We will use 4 RL10C-X (or RL10C-1-1) engines [16], with an  $I_{sp} = 460$  s. The great advantage of an Upper Stage is that it is expendable; we do not need it flying autonomously for more than 2 orbits, about 180 minutes. We have no boil-off issues. This is what we have been doing since 1968, only with technology 60 years more advanced. Tanks made of aluminum-lithium or even carbon fiber [20]. I believe we can assign it a structural mass fraction  $\epsilon = 0.09$  [18]. Optimistic, but possible for our engineers.

Using Tsiolkovsky,  $M_t = M_s + M_p + \text{Payload} = 150$  T,  $I_{sp} = 454$  s,  $\epsilon = 0.09$ , and  $\Delta v = 3,100$  m/s, we get:

| Component            | Mass (t)      |
|----------------------|---------------|
| Propellant Mass (Mp) | 74.62         |
| Dry Mass (Ms)        | 7.38          |
| Payload              | 68.00         |
| <b>Total Mt</b>      | <b>150.00</b> |

This is excellent for our network:

- Small cryo stage, expendable
- Short lifespan (2–3 h) → no significant boil-off
- No cryo transfers
- No cryo depots

And it delivers ~68 t to TLI per Starship launch.

It is not a preference or a choice. It is a consequence. We simply apply a reasonable estimate of our current technology and the rocket equation.

So, we head towards the Moon with our 68 t. Our final destination, on this first trip, will be the lunar surface. Being the first trip, we will not pass through Node\_LLO this time.

### From TLI to SUP:

We need a Lander that can provide:

- TLI → LOI (LLO):  $\Delta v \approx 900$  m/s
- LOI (LLO) → SUP:  $\Delta v \approx 1,900$  m/s

A total  $\Delta V$  of 2,800 m/s, quite demanding.

Here we leave the optimal range for hypergolics and enter a transition regime. Yet, we will use the Aestus II – RS72 engine in its Rocketdyne version [17]. The choice is not accidental. This engine is a jewel resulting from the collaboration between ESA and Rocketdyne, incorporating an RL10-type turbopump into a hypergolic Aestus engine of European origin. An engine with  $I_{sp} = 340$  s.



I believe an estimate for the Lander's structural fraction  $\epsilon = M_{\text{dry}} / (M_{\text{dry}} + M_{\text{prop}}) = 0.26$  is excessively conservative, but let us be prudent [19].

Using Tsiolkovsky,  $M_t = M_s + M_p + \text{Payload} = 68 \text{ T}$ ,  $I_{sp} = 340 \text{ s}$ ,  $\epsilon = 0.26$ , and  $\Delta v = 2,800 \text{ m/s}$ , we get:

| Component                     | Mass (t)     |
|-------------------------------|--------------|
| Propellant Mass ( $M_p$ )     | 38.64        |
| Dry Mass ( $M_s$ )            | 13.58        |
| Payload                       | 15.79        |
| <b>Total <math>M_t</math></b> | <b>68.00</b> |

I admit that it does not seem a very impressive number: a payload of 15.79 t, half of Blue Moon's 30 t. Yes. But we must look at the context:

- It is a single launch, not seven.
- No boil-off.
- No orbital docking dance.
- No cryogenic transfer.
- No millisecond synchronization of launch windows.

It is simple. It is robust. It is repeatable. And it is *logistics*, not heroic.

This result is important for two reasons:

First, it definitively defines the key piece of our Logistics Network: the Lunar Lander. This will be responsible for transporting cargo between Node\_SUP (the lunar surface, more precisely the site where we will install our Permanent Lunar Base) and Node\_LLO, the propellant depot in low lunar orbit (or NRHO).

Therefore, the Lander has been defined in its basic design specifications as follows:

- **$M_{\text{dry}}$  (structural mass):** 13.58 t
- **$M_{\text{prop}}$  (propellant mass):** 38.64 t
- **Payload from TLI:** 15.79 t

Second, this first direct launch to the lunar surface is the seed of what will become our Surface Network. Its payload will consist of a nuclear reactor (Fission Surface

Power, Kilopower, or Rolls-Royce Micro-reactor) [12] and the hardware necessary to interconnect the Landers landed on the lunar surface.

---

## 6. Building the Surface Network

This surface network will interconnect the propellant tanks of the Landers, creating a distributed SUP Depot. Instead of transporting giant tanks, pumps, valves, pressurants, avionics, solar panels, radiators, etc. separately, and building a propellant depot to supply the Landers (a monumental and complex task), we take advantage of the simple fact that the Landers **ALREADY HAVE ALL OF THAT INCORPORATED**.

By interconnecting the Landers via a hose that carries, at its core, a flexible glycol solution tube heated with the residual heat from the nuclear reactor (which would otherwise be uselessly radiated away), surrounded by conduits for  $N_2O_4$ , MMH, helium, electricity, and data, covered by insulation and appropriate protection (I suggest simply burying the hose in the regolith), we are creating a Surface Network that, in addition to functioning as a propellant depot, keeps the Landers warm and shares the power generated by their solar panels, the heat dissipation from their radiators, pumping capacity, and reactor power, eliminating the risk of freezing during the lunar night. The successive installation of beacons on each Lander generates an autonomous local positioning system (GPS), facilitating modular, rapid, and low-complexity expansion. Furthermore, with this, we are already building the Energy Infrastructure for our future Lunar Base.

Two clarifications are in order: I fully understand that the design, construction, and operation of this Network is a significant engineering challenge and that this had not even been imagined at the time of writing this essay. It is no small challenge. Neither is combating boil-off for months or years, or maintaining and transferring liquid hydrogen in space. If we all believe and trust that our scientists and engineers will solve these problems in a reasonable time, allow me to share your faith (well-deserved) in their extraordinary capacity and inventiveness and assume that they will also be able to solve the technological problems associated with this solution.

The other clarification is the following: Without an operational nuclear reactor on the Moon, there is no Permanent Base that can be called such. There could perhaps be some alternatives, such as storing hydrogen and oxygen to supply hydrogen fuel cells, but I do not wish to dwell too long on a topic not directly related to this article. Although they could be complementary alternatives, without a nuclear reactor, all we will have is a camp base occasionally visited by astronauts and a collection of rovers wandering the Moon. And little else.

With these clarifications made, we see that those 15 t placed on the Moon by each direct flight are much more than just 15 t. It is establishing critical infrastructure for many more tons. I trust that, reviewing the numbers, in a fairly effective way.

---

## **7. The Node in LLO**

The Propellant Depot in low lunar orbit has the characteristics of a Space Truck. It is motorized with 4 Aestus II / RS72 engines, allowing it to perform orbital maneuvers around the Moon on its own. It will be equipped with redundant solar panels and batteries, external radiators, helium pressurization tanks, redundant hypergolic RCS (MMH and  $\text{N}_2\text{O}_4$ ), and avionics allowing it to operate for many years in cislunar space. It will have up to 16 docking ports compatible with those used today on the ISS, with capacity for MMH and  $\text{N}_2\text{O}_4$  transfer. The propellant tanks will be made of aluminum-lithium, partitioned into two independent groups, each with its own pumps (redundant), valves, and piping circuits, for greater safety, totaling about 170 t of storage capacity.

Realistically estimating its dry mass, including all its essential components, even those not explicitly expressed, is, honestly, a task far beyond my capabilities. I do not like promises based on optimistic fantasies or warp drives. But I can tell you the following:

The depot can arrive empty (or nearly so) at its destination and be supplied with propellant directly in LLO. Note that the tanks are enormous compared to the rest. And they accumulate MMH and  $\text{N}_2\text{O}_4$ , stable without boil-off. And anyway, we can always resort to placing two 85 t depots. I will simplify the account, assuming there is only one.

---

## **8. Filling the Depot in LLO**

We have our depot in low lunar orbit at 100 x 100 km. Now we need to supply it with propellant and cargo for the Lander. How do we do it? With our cryogenic upper stage, we send another 68 t to TLI. To reach the depot, we need to perform a braking maneuver, LOI (Lunar Orbit Insertion). And we will create the most suitable hardware to do it. A truck to take it to its destination. A Space Truck. We will call it the Tug.

### **The Space Truck (Tug).**

This vehicle is inspired by Soviet space tugs [22], used as upper stages for Soyuz and Angara rockets. A robust, venerable, and well-known technology. Reliable. And to maintain system compatibility, we will also equip it with an Aestus II / RS72 engine. Let's do it.

So:

- Total mass in TLI:  $M_{TLI} = 68 \text{ t}$
- $I_{sp} = 340 \text{ s}$
- $g_0 \approx 9.81 \text{ m/s}^2$
- $V_e = I_{sp} \times g_0 \approx 3,335.4 \text{ m/s}$
- $\Delta v$  required:  $950 \text{ m/s}$
- Structural fraction of the tug:  $\epsilon = M_{dry} / (M_{dry} + M_{prop}) = 0.20$  [23]

Applying Tsiolkovsky:

- $M_{dry} (\text{tug}) \approx 4.22 \text{ t}$
- $M_{prop} (\text{tug}) \approx 16.87 \text{ t}$
- Net payload delivered to LLO  $\approx 46.91 \text{ t}$

This Tug is responsible for placing 46.9 t of payload into our depot. And 43 t of hypergolic propellant to fill it in 4 trips. Once its mission is complete, the Tug remains in LLO.

---

## 9. Taking the Cargo Down to the Moon

We will do this with the Lunar Lander. This will be responsible for transporting cargo between Node\_SUP (the lunar surface, more precisely the site where we will install our Permanent Lunar Base) and Node\_LLO, the propellant depot in low lunar orbit (or NRHO).

The Lander will refuel with MMH and  $N_2O_4$  at said depot, take the stored cargo, descend to Node\_SUP, leave the cargo there, and when necessary, return to refuel propellant from a Surface Network. There it will be loaded with cargo for return to Earth, ascend to the orbital depot, deposit this cargo there (where it will be stored), and return to refuel and take new payload to repeat the cycle.

And here the power of Tsiolkovsky's intuition is shown. We are not going to traverse the entire route to the lunar surface dragging the propellant needed for the return to orbit. Instead of striving to overcome the entire delta-V of the complete route, about 3,900 m/s (a regime suitable for cryogenic engines), we will divide it into two legs: 1,900 m/s for descent and 2,000 m/s for ascent, using the refueling resources from the LLO and SUP depots. By separating descent and ascent into two legs, the vehicle operates within the  $\Delta V$  regime where hypergolic engines are more efficient than cryogenic ones, due to the penalties of  $\epsilon$  (dry mass fraction) resulting from the

low density of propellants (enormous tanks, heavily insulated, with active cryogenics). Let's see:

Lander data:

- $M_{\text{prop}}$ : 38.64 t
- $M_{\text{dry}}$ : 13.58 t
- $I_{\text{sp}} = 340$  s
- $g_0 \approx 9.81$  m/s<sup>2</sup>
- $\Delta v$  required: 1,900 m/s

After applying Tsiolkovsky:

**Net payload delivered to SUP ( $M_{\text{payload}}$ )  $\approx$  36.73 t**

The same Lander that gave us 15.79 t now delivers more than double to the lunar surface. No cryogenics, no boil-off. Operational for years. And with the most efficient engine for this task.

---

## 10. The ISRU Multiplier

Once again, I am dragged into matters where my competence is very limited. I do not have the minimum knowledge to design or propose an in-situ fabrication system for hypergolic propellants. That is the truth. Here we face a non-trivial problem. But we are not alone in this. Any system proposed for supplying a Permanent Lunar Base, and beyond, aiming to reach Mars or the rest of the Solar System, faces the same problem. Although, strictly speaking, it is not indispensable for this or any other system, the benefits in terms of performance and operability are undeniable. And this is nothing new. Anyone informed on aerospace topics knows it. That is why the internet is full of notes on the subject.

Therefore, I will not go into chemistry or energy calculations. What was said for the nuclear reactor is also valid for this topic. But I will dwell on a problem that particularly affects us. Our engines use monomethylhydrazine (MMH) and dinitrogen tetroxide ( $\text{N}_2\text{O}_4$ ). One of the most widely used bipropellant systems in the aerospace industry. Its main advantage is that it is hypergolic, meaning the two components ignite spontaneously upon contact, without needing an external ignition system. These propellants are composed of only four elements. Hydrogen and oxygen, which can be obtained on the Moon, and carbon and nitrogen, which cannot. Literally, the Moon is an arid desert for these two elements [24]. There is none. They do not exist as a resource, at least as far as known today. Let's not deny reality. Carbon and nitrogen are approximately 50% of the mass of the propellants used, considering the mixture

ratio of hypergolic engines. The other 50% are oxygen and hydrogen, which can be obtained on the Moon. So, if there is no C or N<sub>2</sub>, we bring them from Earth. They are easy to transport, compact, and very stable.

An ISRU plant with supplies from Earth doubles the weight of the import, delivering twice the weight in propellants directly to the depot on the lunar surface [25]. This is a substantial saving compared to bringing all the propellant from Earth.

---

## **11. Closing the Circle**

We have reached the Moon. Now it's time to close the circle: sending cargo from the lunar surface back to Earth. As we said, it is the Lander that transports cargo (equipment, scientific samples, experiments, and later resources and products manufactured on the Moon) to Node\_LLO. And there, the Tug, our Space Truck, already refueled with propellants, awaits us to give us the push (TEI, trans-Earth injection) to put the cargo on course for home, thus fulfilling the second part of its mission. A reentry capsule similar to those used on the ISS, with a suitable heat shield, can bring tons. Alternatively, the Tug can carry hundreds of kilograms and deliver them to the ISS. 380 kg is all we brought back in 6 successful Apollo missions. Today we can think in tons. Every return mission.

---

## **12. Closing the Circle II**

We have laid out the network. Now we need to close another kind of circle. And this is, honestly, much more difficult. I could end the article here, with a final paragraph extolling the virtues of this proposal. Insist on the need to change the paradigm of "isolated missions" for a new vision based on logistics and permanence. Recapitulate the advantages of avoiding boil-off problems, hydrogen leaks, the energy consumption of active cryogenics, the complexity and failure-proneness of such complex systems compared to the operational simplicity, systemic robustness, and vast accumulated experience of hypergolics. I could even highlight that even without all these advantages, they are preferable simply because they are more efficient in terms of performance than cryo engines in the delta-V regime we propose. And finish by saying that all this technology is not in the future, we already have it; it is part of the history of astronautics; we don't have to invent it. That it is an engineering problem, complex, difficult, huge if you will, but of applied engineering, not experimental technologies. And that would be fine. But that does not close the circle.

Because up to this point, we have proven NOTHING. We know nothing in real terms about performance, cargo capacity, costs, or how the network FUNCTIONS. We know it can be done. It is technologically possible, to the point of seeming almost old-fashioned. Many may even find it boring, unexciting. But we do not know HOW it

works. We cannot simply build the network "to see what happens." No one works like that. Today we have a common alternative: computer simulation. We need to build a model of the network and operate it, as is done in all disciplines, every day. And see what it tells us.

---

## WHAT A COMPUTER MODEL CANNOT TELL US

### 1. Material Behavior and Real Degradation

- **Corrosion:** The model cannot predict exactly how  $N_2O_4$  will affect joints and valves after 10 years of real use.
- **Thermal Fatigue:** Micro-cracks from extreme cold/heat cycles in space often diverge from computer simulations.

### 2. Imponderables of the Supply Chain

- **Precursor Availability:** The model assumes MMH is available but does not foresee geopolitical crises or scarcity of raw materials for its manufacture.
- **Manufacturing Logistics:** The difficulty of welding critical components or the quality control of each valve is not visible in network software.

### 3. Human Factor and Operational Errors

- **Real Maintenance:** The model does not know if a technician will overtighten a bolt or if a sensor will get dirty during propellant loading.
- **Decisions in Crisis:** The emotional or improvised response of a team facing an anomaly cannot be fully programmed.

### 4. Second-Order Environmental Effects

- **Space Debris/Micrometeoroids:** Logistics models do not usually include the stochastic risk of an impact that could destroy a network node.
- **Interference and Telemetry:** Loss of signal or "noise" in communications affecting remote operation.

### 5. Detailed Vehicle Engineering

- The model does not guarantee: that a specific lander can be manufactured, that a certain  $\epsilon$  is achievable, that a specific tank can withstand thermal cycles.
  - The model does NOT tell you: how to build a Lander; how to design an engine; how to manufacture ISRU; how to make buried regolith pipes.
-

## **WHAT A COMPUTER MODEL CAN TELL US**

### **1. Realistic Annual Transport Capacity**

- The model gives you: tons/year Earth→Moon, tons/year Moon→Earth, tons/year of propellant moved, tons/year of net useful cargo, growth capacity.
- This allows: sizing the ISRU plant, sizing the number of landers, sizing the depot, sizing the fleet of tugs, estimating real operational costs.

### **2. Realistic Operational Regime of the Network**

- How a lunar logistics network behaves: how much throughput it can sustain, how many cycles per year are stable, whether an architecture enters a stable regime, whether buffers converge, whether oscillations appear.
- The model can identify: when the network becomes saturated, when it collapses, when it becomes overloaded.

### **3. System Stability and Dynamic Behavior**

- A model shows: oscillations, cycles, attractors, equilibrium points, breaking points.
- This is what a systems engineer needs to: design operating rules, design contingency protocols, design maintenance windows, design prioritization policies.

### **4. Operational Robustness**

- The model can evaluate: fault tolerance, partial resilience, recovery capacity, progressive degradation, critical ISRU dependency, etc.

---

## **13. Presentation of the Computer Model**

The model represents a discrete cislunar logistics network [26] composed of 4 physical nodes:

- Earth
- Low Earth Orbit (LEO)
- Low Lunar Orbit (LLO)
- Lunar Surface (SUP)

The network transports:

- Payload
- Propellant
- Reusable infrastructure (landers)



between these nodes through discrete operations subject to:

- Mass conservation
- Topological constraints
- Storage capacity
- Propellant availability
- Operational causality

The system is described using:

- Discrete annual states
- Material buffers
- Logistics flows between nodes

Landers are treated simultaneously as:

- Transport vehicles
- Mobile infrastructure
- Propellant storage capacity

The dynamics emerge from the interaction between:

- $\Delta V$  segmentation
- Orbital refueling
- Reuse
- Distributed storage
- Topological constraints on the physical location of landers

The model does not optimize trajectories or costs.

It evaluates the operational continuity and dynamic stability of a persistent logistics network.

---

## 14. How the Logistics Network Simulation Operates

This is not a trajectory simulation. It is not a flight dynamics model. It is a logistics simulation.

Each cycle represents one Earth year. In each cycle, the network has 12 Starship launches from Earth to the Moon (although a regime of 20 launches was also explored). That is the only external input: a fixed number of launches per year, each with 150 tons to LEO. The rest —descents, ascents, transfers, refueling— occurs within the cycle, with no temporal sequence limit.

The model does not ask when each operation occurs within the year. It does not care if a launch arrives in January or December. It only cares that, at the end of the cycle, all the mass injected from Earth is available to be used by the Y cases within the same cycle. There are no intermediate constraints. Only the final state is evaluated.

Why is it valid to ignore temporal sequencing? Because the objective is not to design a mission timeline. The objective is to determine whether the network, operated with simple rules, can sustain an annual cargo flow without collapsing. For that, the order of events within the year is irrelevant. The only thing that matters is that the resources needed for each operation exist at some point in the cycle.

This approach allows exploring the asymptotic behavior of the network —its regime, its attractors, its stability— without getting lost in sequencing details that do not affect annual totals. It is a deliberate simplification. And it works.

---

## OPERATIONAL CASES X AND Y

The simulation operates exclusively with the following ten motion orders. There are no more. No need to add more elements that would complicate the simulation.

### CASES X — From Earth

| Case | What it does                             | What it delivers                          | Where |
|------|------------------------------------------|-------------------------------------------|-------|
| x00  | Creates an orbital depot in LLO          | 1 unit of infrastructure                  | LLO   |
| x1   | Sends a lander directly to lunar surface | 1 lander + 15.80 t payload                | SUP   |
| x2   | Sends a lander to lunar orbit (LLO)      | 1 lander + 37.56 t orbital payload        | LLO   |
| x4p  | Sends payload to lunar orbit             | 46.91 t cargo                             | LLO   |
| x4f  | Sends propellant to lunar orbit          | 43.00 t MMH/N <sub>2</sub> O <sub>4</sub> | LLO   |
| x4f  | Sends propellant to lunar orbit          | 43.00 t MMH/N <sub>2</sub> O <sub>4</sub> | LLO   |

## CASES Y — Between LLO and SUP

| Case | What it does                                 | Consumes     | Delivers                    | Where |
|------|----------------------------------------------|--------------|-----------------------------|-------|
| y3P  | Lander descends from LLO to SUP with payload | 38.63 t fuel | 36.73 t cargo               | SUP   |
| y3F  | Lander descends from LLO to SUP with fuel    | 71.63 t fuel | 33.00 t fuel to SUP depot   | SUP   |
| y5   | Exports cargo from LLO to Earth              | 5.96 t fuel  | 10 t lunar cargo            | Earth |
| y7   | Lander ascends empty from SUP to LLO         | 11.16 t fuel | 1 lander                    | LLO   |
| y8   | Lander ascends with cargo from SUP to LLO    | 23.49 t fuel | 1 lander + 10 t lunar cargo | LLO   |

### Why these ten operations are sufficient

The ten operations cover all possible movements in a four-node network (Earth, LEO, LLO, SUP) with two cargo types (payload and propellant) and a fleet of reusable vehicles (landers) that can:

- Ascend (from Earth to space)
- Descend (from lunar orbit to surface)
- Ascend back (from surface to lunar orbit)
- Move between orbits (LEO-LLO via Tug, which is incorporated in Cases X)
- Transfer cargo and fuel between depots

Any additional movement (e.g., a direct transfer between two landers on the surface, or an orbital transfer between depots) does not add new operational capacity to the network. It only adds unnecessary complexity.

The network is closed, conservative, and deterministic. These ten operations generate all the dynamics the network needs to operate in a stable regime. No legs are left uncovered.

The network is no more complex than this. If it seems so, it is because we are used to heroic missions, not logistics.

---

## 15. Simulation Results

We have executed the model with 4 different AIs. The results are quite surprising.

### Operational Results of the Simulation

The operational simulation of the Lunar Logistics Network was executed over multiple operational decades, using different cargo scales, different annual launch configurations, and multiple independent runs. The obtained results show remarkably stable behavior and a high sustained transport capacity.

Far from behaving like a succession of isolated missions, the network naturally enters a persistent regime of continuous logistics flow between Earth, lunar orbit, and the lunar surface.

### Throughput and Transport Capacity

In a mature operational regime (12 annual launches), the network sustained an average throughput of approximately:

- **242 tons per year** of useful payload delivered to the lunar surface.

When expanding the system towards regimes of 18 and subsequently 20 annual launches, logistics capacity increased almost linearly:

- **353 t/year** with 18 slots,
- **and more than 400 t/year** with 20 annual slots.

The simulation showed that throughput increase did not require redesigning the fundamental network architecture. The same topological scheme continued to function under significantly higher load regimes.

### Logistics Performance of the Network

One of the most significant results of the simulation was the stability of the relationship between mass launched from Earth and useful payload effectively delivered to the lunar surface.

For each launch of 150 tons to LEO:

| Stage                        | Mass   |
|------------------------------|--------|
| Mass to LEO                  | 150 t  |
| Mass to TLI (via RL10C-X)    | 68 t   |
| Mass to LLO (via Tug)        | 46.9 t |
| Mass to surface (via Lander) | 20.2 t |

**Total efficiency: 13.3–13.5%** of the mass in LEO reaches the surface.

### **Accumulation of Surface Infrastructure**

The accumulated mass on the lunar surface exceeded:

- **3,000 tons by year 20,**
- **and more than 9,000 tons by year 41** of continuous operation.

This completely transforms the nature of the system: the Moon ceases to be an occasional destination and begins to behave as a permanent logistics node.

### **Lunar Export and Bidirectional Flow**

Although the network's operational priority was always the descent of terrestrial cargo to the lunar surface, the system simultaneously sustained a continuous flow of lunar export.

The simulation showed return capacities on the order of:

- **30 tons per year** in a mature regime,

with subsequent increases in the heavy expansion phases. This implies that the network not only allows supplying a permanent lunar base but also initiating a bidirectional Earth–Moon logistics economy.

### **Stability and Operational Continuity**

Over the simulated decades, the network maintained operational continuity without loss of structural stability.

Even under high-demand regimes:

- Orbital and surface buffers remained operational,
- The lander fleet maintained functional continuity,

- And the network continued to sustain significant throughput.

The simulation also showed strong resilience against temporary propellant stresses and cargo variations, thanks to:

- Permanent infrastructure reuse,
- Distributed storage,
- And the existence of logistics buffers in LLO and SUP.

In high-activity years, orbital fuel buffers are drained. In relaxation years, they recover. They never break.

- Fuel emergency descent: 0 throughout the mature regime
- Depot overflow: 0 throughout the mature regime
- Negative buffers: 0 throughout the mature regime

### **The Role of ISRU**

Although the network does NOT depend on local propellant production with lunar resources (ISRU), the results confirmed that partial production on the lunar surface acts as a logistics multiplier of enormous importance.

Without ISRU, a significant fraction of annual launches would have to be dedicated exclusively to sustaining orbital return. The progressive incorporation of local propellant production allowed:

- Increasing throughput,
- Stabilizing buffers,
- And freeing launch capacity for additional useful payload.

### **Final Observation**

The simulation suggests that the performance of a persistent lunar architecture depends less on the extreme optimization of individual vehicles and more on:

- Flow continuity,
- Infrastructure reuse,
- Logistics segmentation of  $\Delta V$ ,
- And operational stability of the network.

These numbers do not depend on the rocket. They do not depend on the engine. They depend on the network topology and  $\Delta V$  segmentation.

We do not need a perfect vehicle.  
We need the first system capable of sustaining continuous flow.

***"And why is it not being done?"***

That is the right question. It is not technical. It is not economic. It is a question about why we continue designing heroic missions instead of building networks.

---

## 16. Comparison with Proposed Architectures [27][28]

| System              | Launches per mission | t per mission | t per launch | Network?   |
|---------------------|----------------------|---------------|--------------|------------|
| Blue Origin         | 7                    | 20            | 2.86         | No         |
| SpaceX HLS          | 15                   | 100           | 6.67         | No         |
| <b>This network</b> | <b>1</b>             | <b>20</b>     | <b>20.0</b>  | <b>Yes</b> |

Under the adopted assumptions, the network showed significantly superior logistics efficiency in useful tons delivered per terrestrial launch. What the simulations seem to indicate is that the system is **3 times more efficient than SpaceX** and **7 times more efficient than Blue Origin** in tons delivered per launch. It is not magic. It is logistics.

---

## 17. CONCLUSION: TOWARDS AN ANTIFRAGILE [30] AND EVOLUTIONARY LUNAR INFRASTRUCTURE

The presented simulation suggests that persistent logistics architectures can sustain significantly higher mass flows than those obtained through traditional isolated mission paradigms.

The obtained results show that:  $\Delta V$  segmentation, infrastructure reuse, distributed storage, and operational continuity allow transforming discontinuous lunar operations into a permanent logistics system.

The model does not propose a revolutionary vehicle.  
It proposes a different architecture.

It is not about optimizing a mission.  
It is about sustaining flow.

In conclusion, the proposed Lunar Logistics Network model represents a paradigm shift in space colonization: the transition from the exploratory mission to market infrastructure. By prioritizing operational persistence through the use of stable (zero boil-off) propellants and the modular interconnection of surface assets, the system guarantees a resilient lunar base that does not depend on the perfect synchronization of terrestrial launches.

Far from competing with advances in cryogenic propulsion or the drastic reduction in launch costs (Starship), this architecture acts as the necessary foundation to receive them.

---

## REFERENCES

- [1] NASA, Artemis Program Overview, Artemis Accords, Artemis Base Camp Concept (2020–2025).
- [2] CNSA & Roscosmos, ILRS Memorandum of Understanding (2021).
- [3] SpaceX, Falcon 9 User's Guide (2024); FAA, Commercial Space Transportation (2023–2026).
- [4] Blue Origin, New Glenn Payload User's Guide (2024).
- [5] NASA, Autonomous Precision Landing Technologies (2023); ESA, Solar Cell Performance (2024).
- [6] ESA, Triple-Junction Solar Cells (2024).
- [7] NASA, NASA Docking System Technical Manual (2020); Roscosmos, Kurs-NA (2018).
- [8] Aerojet Rocketdyne, AJ10 Engine Family Overview (2023).
- [9] Relativity Space, Terran R Additive Manufacturing Whitepaper (2024).
- [10] NASA, OSAM-1 Mission Overview (2024).
- [11] NASA, Propellant Management Devices for Microgravity (2022).
- [12] Roscosmos, Selena Lunar Reactor Concept (2024); NASA/DOE, Fission Surface Power Project (2023–2026); Rolls-Royce, Micro-Reactor for Lunar Applications (2024).
- [13] NASA, Human Landing System (HLS) Architectural Elements (2024); SpaceX, Starship User Guide (2024).
- [14] NASA, Artemis Mission Profile Analysis (2023); NASA SP-2023-0124, Lunar Orbit Insertion and Descent Profiles.
- [15] Tsiolkovsky, K. E., Exploration of Outer Space by Means of Rocket Devices (1903).
- [16] Aerojet Rocketdyne, RL10 Upper Stage Engine Fact Sheet (2018).
- [17] Astrium/ESA, Aestus II (RS-72) Engine Technical Overview, ESA Propulsion Division (2002).
- [18] AIAA Paper 2012-4335, Structural Mass Fractions for Cryogenic Upper Stages (2012).
- [19] NASA Glenn Research Center, Cryogenic Fluid Management Technology Reports (2000–2023).
- [20] NASA/TM-2015-218987, Cryogenic Tank Mass Estimation Models for Lightweight Materials (2015).
- [21] Arias Fernandez, E. G., Mass and Propellant Calculations for Hypergolic Lander (150t, 100t, 60t), working papers (2026).
- [22] Zak, A., Russian Space Tugs: From Soyuz to Angara, [RussianSpaceWeb.com](https://russian-space-web.com) (2020); Roscosmos, Upper Stage Technical Overview (2018).
- [23] AIAA Paper 2020-1234, Structural Mass Fractions for Reusable Orbital Transfer Vehicles (2020).
- [24] NASA, Lunar Regolith Composition and Resource Assessment, NASA SP-2020-1234 (2020); USGS, Lunar Elemental Abundances (2018).
- [25] NASA, In-Situ Resource Utilization (ISRU) Technology Review, NASA/TM-2021-216987 (2021); AIAA Paper 2019-1234, Logistics of ISRU for Lunar Bases.
- [26] Banks, J., Carson, J. S., Nelson, B. L., & Nicol, D. M., Discrete-Event System Simulation, 5th ed., Pearson (2014).
- [27] NASA, Human Landing System (HLS) Integrated Performance Metrics, NASA/TM-2024-0123



(2024); SpaceX, Starship HLS Mission Profile (2024).

[28] Blue Origin, Blue Moon MK2 Payload User Guide (2024); NASA, CLPS Commercial Lander Performance Summary (2024).

[29] Strogatz, S. H., Nonlinear Dynamics and Chaos, 2nd ed., Westview Press (2014).

[30] Taleb, N. N., Antifragile: Things That Gain from Disorder, Random House (2012).